

Instruction-set Selection for Multi-application based ASIP Design: An Instruction-level Study

Roshan Ragel*, Swarnalatha Radhakrishnan*, and Angelo Ambrose†

* Department of Computer Engineering
University of Peradeniya, Sri Lanka

† School of Computer Science and Engineering
University of New South Wales, Australia

Abstract—Efficiency in embedded systems is paramount to achieve high performance while consuming less area and power. Processors in embedded systems have to be designed carefully to achieve such design constraints. Application Specific Instruction set Processors (ASIPs) exploit the nature of applications to design an optimal instruction set. Despite being not general to execute any application, ASIPs are highly preferred in the embedded systems industry where the devices are produced to satisfy a certain type of application domain/s (either intra-domain or inter-domain). Typically, ASIPs are designed from a base-processor and functionalities are added for applications. This paper studies the multi-application ASIPs and their instruction sets, extensively analysing the instructions for inter-domain and intra-domain designs. Metrics analysed are the reusable instructions and the extra cost to add a certain application. A wide range of applications from various application benchmarks (*MiBench*, *MediaBench* and *SPEC2006*) and domains are analysed for two different architectures (*ARM-Thumb* and *PISA*). Our study shows that the intra-domain applications contain larger number of common instructions, whereas the inter-domain applications have very less common instructions, regardless of the architecture (and therefore the ISA).

I. INTRODUCTION

Embedded systems are the realm in current civilisation and their omni-presence in modern technology in the form of mobile phones, network devices, computers, medical devices, automotive and other applications is obvious. Power and energy consumption, device size, durability and reliability are some of the major properties which are expected from such embedded devices. Hence it is imperative that the embedded systems be optimised for the needs of its application to achieve maximum efficiency. Application Specific Integrated Circuits (ASICs) are specifically made in hardware to execute a functionality with extremely efficient power, area and performance budgets. Despite being heavily used in the industry for System-in-Chip designs, ASICs are hardly flexible and can not be reused for a different type of application. Field Programmable Gate Arrays (FPGA), on the other hand, are highly flexible, but inefficient for power and performance. FPGAs are still considered as prototyping platforms for embedded systems mainly due to their inefficiency in area and power. Figure 1 depicts an illustrative diagram of different technologies including ASIC and FPGA. As shown, the ASIC is efficient but lacks flexibility, whereas the FPGAs are flexible, nevertheless costs performance, power and area. To hit a reasonable balance between ASICs and FPGAs,

Application Specific Instruction-set Processors (ASIPs) are considered as the appropriate choice. As shown in Figure 1, ASIPs, which are the latest technological trend in embedded systems [6], are conceived by tightening up the flexibility from FPGAs and releasing the efficiency from ASICs.




Fig. 1. ASIC vs. ASIP vs. FPGA

ASIPs are formed using a hardware/software co-design process, where instructions are chosen for a processor, based on the behaviour of the application to be deployed. Applications are designed using instructions which are then executed in a processor. Such a technology improves the design productivity due to the simplicity in the software implementation process. Furthermore, the hardware/software co-design approach improves productivity by allowing the hardware to be reused and reprogrammed [6]. The complexity of the design is reduced, thus Non-Recurring Engineering (NRE) cost is also decreased. The instruction set of an ASIP is tailored to benefit a specific application or a known set of applications. Such an instruction set based solution provides high degree of flexibility, supporting yet-to-be introduced standards and implementations. This provides reasonable tolerance in design changes which might arise in the future. ASIPs are typically modelled using high level languages [6], which allows a relatively easy and methodical approach to design applications on a resource stringent hardware.

Instructions in an ASIP is an integral component to decide the functionality and its efficiency. Since complex application programs use hundreds of types of processor instructions, selecting and designing most suitable instructions to achieve the highest performance in an optimised way is a major

challenge in the design process of an ASIP. An instruction set (also known as the instruction set architecture, ISA) serves as the interface between hardware and software in a computer system. In an application specific environment, the system performance can be improved by designing an instruction set that matches the characteristics of the hardware and the application [11].

From a cost and performance perspective, types of instructions used in given applications is vital. Approaches to instruction set generation for an ASIP can be classified as either instruction set synthesis approach [7, 11] or instruction set selection approach [3, 12, 16, 19] on the basis of how the instructions are generated. In synthesis approach, instruction set is synthesised for a particular application based on the application requirements, while in selection approach, a superset of instructions is available and a subset of them is selected to satisfy the performance requirements within the architectural constraints [13].

Instructions in an application are affected by three factors: 1), functionality of the application and its relationship to the ISA; 2), behaviour of the compiler's code generation; and, 3), coding style. We only focus on the first one which is the most critical of all. In this paper we perform an instruction-level study to realise the nature of instructions used within application domains (i.e., intra-domain) and across application domains (i.e., inter-domain). Such a study allows us to envision the effect on instruction commonalities and uniqueness for application specific instruction sets. This paper provides an insight into the instruction usage in applications to evaluate the intra-domain and inter-domain costs involved for integration.

II. MOTIVATION

The efficacy of ASIP applications depends on the optimal use of the instructions. Figure 2 illustrates three different application domains; automotive, multimedia and security. The Security domain is illustrated with three applications: AES, DES and RSA, combined in intra-domain. If we are to design an embedded system to include applications from these three domains (i.e., inter-domain), it is necessary to realise the extra cost involved for integration in terms of instructions which is directly related to the design time and effort. Since the application domains have quite a significant functional difference, we expect to find very less commonality (very less reusable instructions) in the instructions across the three domains. The applications inside an application domain (i.e., intra-domain) are expected to contain much less uniqueness in instructions (less additional cost and high reusability¹) across different applications, due to similar type of operations (i.e., functionalities) being performed. We endeavour to validate this hypothesis by studying the instructions being used in inter-domain and intra-domain applications. It is further important to evaluate the contribution of the instruction set to the ASIP design, compared to the coding style and the nature of the compiler.

¹we refer to this intersecting instructions as reusable instructions which are typically built as base processors in state-of-the-art ASIPs [1]




Fig. 2. Application Domains, Applications and Instructions

Embedded systems are evolving rapidly and the amount of applications executed in an embedded system (such as mobile phones) range in the order of tens to hundreds, while still growing. Hence it is necessary for the designer to realise the additional cost involved in integrating applications and the means of reusability to improve the design process. We evaluate these two properties: 1), extra cost and 2), reusability, at the instruction level.

III. RELATED WORK

ASIP systems have become the norm in embedded systems to achieve high performance while being able to consume low power [2, 6]. Selection of instructions for an ASIP has been widely studied. We discuss the most appropriate studies in this section.

One approach for generating instruction sets is by considering the datapath model. In 1994, Praet et al. [20] have shown how instruction selection for ASIPs can be performed by generating a combined instruction set and datapath model from the instruction set. Operation bundling was performed on the model with an abstract datapath. This methodology still requires refinement and testing. Then Kucukcakar [14] came up with an architecture and a co-design methodology to improve the performance of embedded system applications through instruction-set customisation, based on a similar kind of concept. Although these methodologies improve the performance of ASIPs they failed to consider the design constraints such as area, power consumption, NRE cost etc.

One of the early work on methodologies to maximise performance of ASIP under design constraints, such as area, power consumption and NRE cost, is [4]. The authors in [4] proposed a rapid instruction selection approach from their library of pre-designed specific instructions to be mapped on a set of pre-fabricated co-processors/functional units. As a result, the authors in [4] were able to significantly increase application performance while satisfying area constraints. This methodology uses a combination of simulation, estimation and a pre-characterised library of instructions, to select the appropriate co-processors and instructions. Alomary et al. [3] proposed a new formalisation and an algorithm that considers the functional module sharing. This method allows designers to predict the performance of their designs before implementation, which is an important feature for producing a high quality design in reasonable time. In addition to that, an efficient algorithm for automatic selection of new application-specific instructions under hardware resources

TABLE I
BENCHMARK APPLICATIONS USED

Applications per Domains					
Automotive (AM)	Office (OF)	Security (SE)	Telecomm (TC)	MediaBench (MB)	Spec.CPU2006Int (SP)
BasicMath	GhostScript	Blowfish	Adpcm	Epic	BZip2
BitCount	ISpell	PGP	CRC32	G721	MCF
QuickSort	RSynth	Rijndael	FFT	H263enc	Hammer
Susan	StringSearch	Sha	GSM	MPEG2enc	Sjeng

constraints is introduced in [5]. The main drawback of this algorithm is the un-optimised Very-High-Speed Integrated circuits Hardware Description Language (VHDL) model.

Researchers have already proposed automated techniques in ASIP design process to achieve best performance under certain design constraints. Almer et al. [17] presented a complete tool-chain for automated instruction set extension, micro-architecture optimisation and complex instruction selection, based on GCC compiler. Huang and Despain in [10] proposed a single formulation, combining the problem of instruction set design, micro-architecture design and instruction set mapping. The formulation receives as inputs the application, architecture template, objective function and design constraints, and generates as outputs the instruction set for the application. Similarly, Zhu et al. [22] presented a design automation approach, referred to as Automatic Synthesis of Instruction-set Architectures (ASIA), to synthesise instruction sets from application benchmarks. The problem of designing instruction sets was formulated as a modified scheduling problem in [22]. In [21], a design flow was proposed to automatically generate Application-Specific Instructions (ASIs) to improve performance with memory access considerations. The ASIs are selected not only based on the instruction latency but also the memory access.

Once the instructions are chosen for an ASIP, the selected instructions are evaluated. Authors in [11] and [16] introduced methods to evaluate instruction sets with several design constraints. Peymandoust et al. [18] automatically grouped and evaluated data-flow operations in the application as potential custom instructions. A symbolic algebra approach is utilised to generate the custom instructions with high level arithmetic optimisations.

Considering the process of instruction selection and evaluation for ASIPs, in this paper, we perform an application analysis (at the instruction level) to identify the commonalities and uniqueness for intra-domain and inter-domain applications. We evaluate the applications based on the extra cost for application integration and reusability of common instructions. Such an analysis will enlighten the designer in performing smart instruction selection.

IV. METHODOLOGY

As highlighted in Section II, our objective is to study the reusability and extra cost of multi-application ASIPs (named mASIPs) in terms of instruction set utilisation. The method we device to perform this study is described in this section. For every instance of our experiment, we choose a set of target applications, one or many of which can be deployed in our mASIP. Therefore, the target application set is a list of

potential applications for an mASIP design. The target applications can come from a single application domain (such applications are identified as intra-domain applications) or multiple application domains (such applications are identified as inter-domain applications).

With respect to the instruction set design, the mASIPs can be built using two phases: one, designing and building a *base processor* with the instruction set necessary for all the applications of the target set and two, extending the base processor to cater the rest of the instruction types for the applications to be deployed for a particular mASIP. It is worth to note that the applications that will be deployed for a particular mASIP is a subset of the applications in the target set. We will define the instruction set of the base processor as the *base instruction set*. Therefore, the *base instruction set* is the set of instructions that are common to all the applications in our target application set.

With this background, in our study, we calculate the reusability and the extra cost of an mASIP by using the *base instruction set* and the rest of the instruction set necessary for building an mASIP. That is, a larger base instruction set will indicate a higher reusability and a larger additional instructions in phase two of our design would indicate a higher extra cost. Both reusability and extra cost of a particular mASIP design will be quantified by the number of instructions in the base instruction set and the rest of the instructions needed to complete the mASIP design.




Fig. 3. Experimental Flow

V. EXPERIMENTAL SETUP

Figure 3 explains the experimental flow of our study. Application benchmarks are assembled, using two different cross compilers (targeting two well known instruction set architectures, *ARM-Thumb* and *PISA*), to create the assembly files, indicating all the instructions used. It is worth to note that *ARM-Thumb* has 78 instruction in its ISA and *PISA* has 72 instructions (integer instructions only) in its ISA. We call them the complete instruction sets for *ARM-Thumb* and *PISA*. Applications are collected from different domains, four applications each. The assembly files are then analysed for intra-domain and inter-domain instruction level

$$Reusability\ Factor = \frac{(\# of\ instructions)_{Base\ ISA} * 100}{(\# of\ instructions)_{mASIP}} \quad (1)$$

$$Extra\ Cost\ Factor = \frac{[(\# of\ instructions)_{Apps/Domain} - (\# of\ instructions)_{Base\ ISA}] * 100}{(\# of\ instructions)_{mASIP}} \quad (2)$$

dependencies. The complete set of ISA of each architecture is another input to the analysis.

TABLE II
INSTRUCTION-SET SELECTION FOR INTRA-DOMAIN APPLICATIONS

Group	Domain	Number of Instructions					
		ARM-Thumb ISA			PISA		
		Indiv.	Inter.	Union	Indiv.	Inter.	Union
AM	BasicMath	33			25		
	BitCount	46	23	49	31	16	40
	QSort	25			19		
	Susan	45			34		
OF	GhostScript	52			44		
	ISpell	29	27	55	50	27	51
	RSynth	52			40		
	StringSearch	40			27		
SE	BlowFish	49			30		
	PGP	57	30	57	52	22	52
	Rijndael	36			30		
	Sha	40			29		
TC	Adpcm	39			32		
	CRC32	36	25	55	22	16	45
	FFT	41			30		
	GSM	54			41		
MB	Epic	56			44		
	G721	49	45	56	41	35	50
	MPEG2	51			43		
	Rasta	53			44		
SP	BZip2	57			50		
	Hammer	45	45	58	29	29	52
	Sjeng	55			46		
	H264	54			47		

Table I lists the applications used in our study from different benchmarks. We have identified the applications under six domains. Four of the six domains are coming from the famous *MiBench* benchmark suite [8] and they are *Automotive* (AM), *Office* (OF), *Security* (SE) and *Telecomm* (TC). The next domain contains four applications from *MediaBench* [15] benchmark suite and the last domain is a set of integer applications from *Spec2006* CPU [9] benchmark suite.

VI. RESULTS AND ANALYSIS

In this section we present the results we obtained from our extensive instruction-level study of reusability and extra cost of the mASIPs with a carefully selected set of target applications. Reusability of instructions, the *Reusability Factor*, is defined as in Equation 1 using the *base instructions* and extra cost of supporting an application/domain, the *Extra Cost Factor*, is defined using Equation 2. We further analyse the results in order to identify suitable patterns and behaviours that could be used in building a multi-application based ASIP design automation tool.

In Table II, we show the instruction set selection for intra-domain target application set. That is, a particular target mASIP can only be deployed with applications from a single

domain. Therefore, we have repeated the experiment six times, one for each domain targeted and the results are reported. It is worth to note that the instruction set selection for these experiments are groups separately for *ARM-Thumb* (columns 3-5) and *PISA* (columns 6-8). Columns 3,4, and 5 gives the number of instructions of individual application, intersection of all four applications of the particular domain, and the union of four applications of the same domain in case of *ARM-Thumb* ISA. Similar results are reported for *PISA* in columns 6,7, and 8.

Let us take one of these six experiments of the *ARM-Thumb* ISA, *Automotive* of Table II, which contains the following applications: *BasicMath*, *BitCount*, *QSort* and *Susan*. For this experiment, the target applications are the four mentioned earlier and therefore our mASIP can support one or many of the four applications. Therefore, we have computed the intersection of the instruction sets from these four applications as our base instruction set and this number is 23. Now, if we are to deploy *BasicMath* (which is having a total of 33 instructions as shown in column 3 of Table II) on top of our base instruction set, we need to include 10 more instructions. Similar numbers for *BitCount*, *QSort* and *Susan* are 23, 2 and 22 respectively. In addition, if we are to deploy all four applications at the same time, the total number of instructions required are 49 including the base instruction set, this is given in column 5 as the union value. The rest of the figures in Table II are similar results for the experiments conducted in the rest of the domains namely *Office*, *Security*, *Telecomm*, *MediaBench* and *Spec2006*.

From the values in Table II, the number of *base instructions* as a percentage to the union, total number of instructions in the mASIP (the *Reusability Factor* as per Equation 1) are calculated and are: 47%, 49%, 53%, 45%, 80% and 78% for *Automotive*, *Office*, *Security*, *Telecomm*, *MediaBench* and *Spec2006* respectively. The average (arithmetic mean) of these numbers is 59% and can be considered as our mean *Reusability Factor* for the six experiments we conducted for *ARM-Thumb*. Using the values in the same table, the numbers of instructions required to deploy a particular application on a mASIP in addition to the *base instruction* set of the domain as a percentage to the union (the *Extra Cost Factor* as per Equation 2) are calculated and the average values for each domain are: 29%, 30%, 27%, 32%, 13%, and 13% for *Automotive*, *Office*, *Security*, *Telecomm*, *MediaBench* and *Spec2006* respectively. The average (arithmetic mean) of these numbers is 24% and can be considered as the mean *Extra Cost Factor* for the six domains for *ARM-Thumb*. Given that the experiments are for intra-domain applications, the reusability is

TABLE III
INSTRUCTION-SET SELECTION FOR INTER-DOMAIN APPLICATIONS

Group	Domain	Number of Instructions					
		ARM-Thumb ISA			PISA		
		Indiv.	Inter.	Union	Indiv.	Inter.	Union
SET-01	AM	45			37		
	OF	51	14	54	48	12	48
	MB	53			46		
	SE	54			48		
SET-02	AM	45			37		
	OF	51	15	55	48	13	48
	MB	53			46		
	SP	55			48		
SET-03	AM	45			37		
	OF	51	14	53	48	11	48
	MB	53			46		
	TC	52			42		
SET-04	AM	45			37		
	OF	51	14	55	48	12	48
	SE	54			48		
	SP	55			48		
SET-05	AM	45			37		
	OF	51	13	54	48	11	48
	SE	54			48		
	TC	52			42		
SET-06	AM	45			37		
	OF	51	14	55	48	11	48
	SP	55			48		
	TC	52			42		
SET-07	AM	45			37		
	MB	53	16	55	46	12	48
	SE	54			48		
	SP	55			48		
SET-08	AM	45			37		
	MB	53	14	54	46	11	48
	SE	54			48		
	TC	52			42		
SET-09	AM	45			37		
	MB	53	17	55	46	11	48
	SP	55			48		
	TC	52			42		
SET-10	AM	45			37		
	SE	54	14	55	48	11	48
	SP	55			48		
	TC	52			42		
SET-11	OF	51			48		
	MB	53	18	55	46	15	48
	SE	54			48		
	SP	55			48		
SET-12	OF	51			48		
	MB	53	16	54	46	13	48
	SE	54			48		
	TC	52			42		
SET-13	OF	51			48		
	MB	53	17	55	46	12	48
	SP	55			48		
	TC	52			42		
SET-14	OF	51			48		
	SE	54	16	55	48	12	48
	SP	55			48		
	TC	52			42		
SET-15	MB	53			46		
	SE	54	19	55	48	13	48
	SP	55			48		
	TC	52			42		

expected to be higher than the extra cost (i.e., applications in an intra-domain should have very similar functionalities hence would require a similar set of instructions) which can be verified from the numbers we obtained here. From Table II we obtained similar results for *PISA*. The mean *Reusability Factor* for the six experiments we conducted for *PISA* is 49% and the mean *Extra Cost Factor* for the six domains for *PISA* is 26% proving our expectation that reusability is higher than the extra cost in intra-domain applications.

In Table III, we tabulate the instruction set selection for inter-domain target application sets. That is, a particular target mASIP can be deployed with applications from different application domains. We assume that a mASIP can integrate at most four application domains (the rest of the combinations are not reported due to lack of space) and given that we have six domains, we could have 15 combinations (6C_4). As shown in Table III, we named each of this combination a SET and therefore we have repeated the experiment 15 times, one for each SET and the results are reported. Columns 3, 4, and 5 show the number of instructions for each domain, the intersection, and the union of all four domains of the particular set respectively in case of *ARM-Thumb* ISA. Columns 6-8 show similar results for *PISA*.

Let us consider the numbers in table III. The domains taken into the set (*SET-01*) are *Automotive* (*AM*), *Office* (*OF*), *MediaBench* (*MB*) and *Security* (*SE*) which contain a total of 16 applications, four each. For this experiment, the target applications are the 16 mentioned earlier and therefore our mASIP can support one or many of the sixteen. Therefore, in case of *ARM-Thumb* ISA, we have computed the intersection of the instruction sets from these 16 applications as our base instruction set and this number is 14. Now, if we are to deploy all the applications in *Automotive* on top of our *base instruction* set, we need to include 31 more instructions as shown in Table III. Similar numbers for *Office* (*OF*), *MediaBench* (*MB*) and *Security* (*SE*) are 37, 39 and 40 respectively. In addition, if we are to deploy all 16 applications at the same time, the total number of instructions required are 54 including the *base instruction* set. The rest of the figures in Table III are for the experiments conducted for the rest of the 14 sets of *ARM-Thumb* and 15 sets of *PISA*.

From the numbers in Table III, in case of inter domain *ARM-Thumb* experiments the mean *Reusability Factor* is 28% and the mean *Extra Cost Factor* is 67%. Given that the experiments are for inter-domain applications, the *Reusability Factor* is expected to be lower than the *Extra Cost Factor* (i.e., applications from different domains will have quite varying functionalities, hence contain different instructions) which is reflected in our experiments.

By using the values shown in table III similar results are calculated in case of *PISA*. The mean *Reusability Factor* is 25% and the mean *Extra Cost Factor* is 68% still proving our expectations that the reusability is lower than the extra cost for the inter-domain scenario, even for a different architecture.

In Figure 4 the four graphs are depicting the reusability factor values and extra cost factor values we have discussed previously in Intra-domain and Inter-domain experiments for *ARM-Thumb* and *PISA* target architectures. In Figure 4 (a), (b) representing intra-domain experiments the mean of reusability factor is higher than the mean of extra cost factor whereas, in Figure 4 (c), (d) representing inter-domain experiments the mean of extra cost factor is higher than the mean of reusability factor.

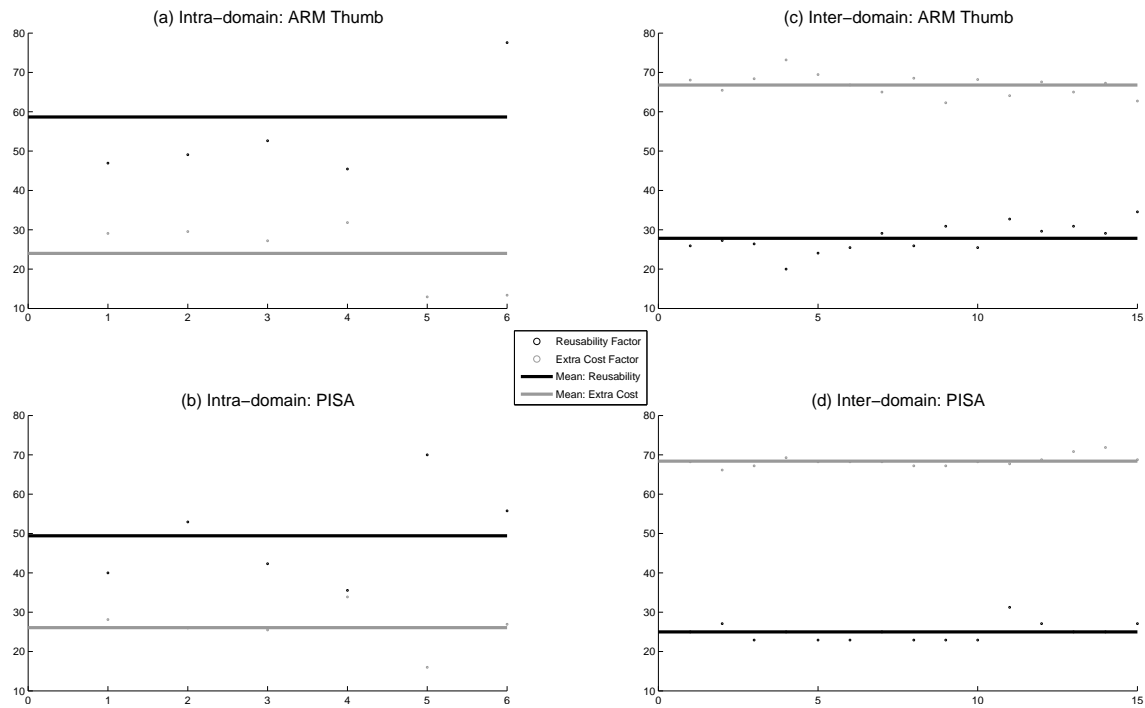


Fig. 4. Reusability and Extra Cost Factors for Intra- and Inter-Domain Applications

VII. CONCLUSION

We performed an extensive study in the instructions for a multi-application based ASIP, which was meant to execute inter-domain and intra-domain applications. *MiBench*, *MediaBench* and *SPEC2006* benchmarks are experimented for *ARM-Thumb* and *PISA* architectures. Our experiments prove that the *reusability* in instructions is larger than the *extra cost* for intra-domain and smaller for inter-domain applications. This justifies our hypothesis that instruction-level analysis is useful to design multi-application based ASIPs, regardless of the instruction set architecture.

REFERENCES

- [1] Xtensa Processor. Tensilica Inc. (<http://www.tensilica.com>).
- [2] Algotochip. <http://www.algotochip.com/>.
- [3] A. Alomary, T. Nakata, Y. Honma, M. Imai, and N. Hikichi. An ASIP instruction set optimization algorithm with functional module sharing constraint. In *ICCAD*, pages 526–532, 1993.
- [4] N. Cheung, J. Henkel, and S. Parameswaran. Rapid configuration and instruction selection for an asip: a case study. In *DATE*, pages 802 – 807, 2003.
- [5] C. Galuzzi, E. M. Panainte, Y. Yankova, K. Bertels, and S. Vassiliadis. Automatic selection of application-specific instruction-set extensions. In *CODES+ISSS*, 2006.
- [6] T. Gikler and H. Meyr. *Design of Energy-Efficient Application-Specific Instruction Set Processors*. Springer Publishing Company, Incorporated, 1st edition, 2010.
- [7] M. Gschwind. Instruction set selection for asip design. In *CODES*, pages 7 – 11, 1999.
- [8] M. R. Guthaus, J. S. Ringenberg, D. Ernst, T. M. Austin, T. Mudge, and R. B. Brown. Mibench: A free, commercially representative embedded benchmark suite. In *Proceedings of the Workload Characterization, 2001. WWC-4. 2001 IEEE International Workshop, WWC '01*, pages 3–14, Washington, DC, USA, 2001. IEEE Computer Society.
- [9] J. L. Henning. Spec cpu2006 benchmark descriptions. *SIGARCH Comput. Archit. News*, 34(4):1–17, Sept. 2006.
- [10] I.-J. Huang and A. M. Despain. Generating instruction sets and microarchitectures from applications. In *ICCAD*, pages 391–396, 1994.
- [11] I.-J. Huang and A. M. Despain. Synthesis of Application Specific Instruction Sets. In *IEEE Trans. in CAD*, 1995.
- [12] M. Imai, N. Binh, and A. Shiomi. A new HW/SW partitioning algorithm for synthesizing the highest performance pipelined ASIPs with multiple identical FUs. In *EURO-DAC*, pages 126–131, 1996.
- [13] M. K. Jain, M. Balakrishnan, and A. Kumar. ASIP Design Methodologies: Survey and Issues. In *VLSID*, page 76, 2001.
- [14] K. Küçükçakar. An ASIP design methodology for embedded systems. In *CODES*, pages 17–21, 1999.
- [15] C. Lee, M. Potkonjak, and W. H. Mangione-Smith. Mediabench: a tool for evaluating and synthesizing multimedia and communications systems. In *Proceedings of the 30th annual ACM/IEEE international symposium on Microarchitecture, MICRO 30*, pages 330–335, Washington, DC, USA, 1997. IEEE Computer Society.
- [16] C. Liem, T. May, and P. Paulin. Instruction-set matching and selection for DSP and ASIP code generation. In *EUROASIC*, pages 31 –37, 1994.
- [17] O. Almer, R.V.Bennett, I.Böhm, A.C.Murray, X.Qu, M.Zuluaga, B.Franke, and N.P.Topham. An End-to-End Design Flow for Automated Instruction Set Extension and Complex Instruction Selection based on GCC. In *GROW*, 2009.
- [18] A. Peymandoust, L. Pozzi, P. Jenne, and G. De Micheli. Automatic instruction set extension and utilization for embedded processors. In *ASAP*, pages 108 – 118, 2003.
- [19] J. Shu, T. C. Wilson, and D. K. Banerji. Instruction-Set Matching and GA-based Selection for Embedded-Processor Code Generation. In *VLSID*, pages 73–78, 1996.
- [20] J. Van Praet, G. Goossens, D. Lanneer, and H. De Man. Instruction set definition and instruction selection for ASIPs. In *ISSS*, pages 11–16, 1994.
- [21] J. Wu, C. Lin, D. Chen, and Y. Wang. Memory models for an application-specific instruction-set processor design flow. *Embedded Software and Systems, Second International Conference on*, pages 471–478, 2008.
- [22] A. Y. Zhu, X. Li, L. T. Yang, and J. Yang. A fast instruction set evaluation method for ASIP designs. In *EUC*, pages 489–498, 2006.